

Re: Why can setuid programs regain root after dropping it when using capabilities?

Re: Why can setuid programs regain root after dropping it when using capabilities?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-12/msg08699.html>

- *From:* daw@xxxxxxxxxxxxxxxxxx (David Wagner)
 - *Date:* Sat, 31 Dec 2005 20:58:50 +0000 (UTC)
-

David Chau wrote:

```
>While debugging some code, I found that a setuid program could regain
>root after dropping root if the program used capabilities. (I tested
>this on 2.6.14 and 2.6.9.) Is this the expected behavior? Here's a
>short test case:
>
> /* chown root this program, suid it, and run it as non-root */
> #include <sys/types.h>
> #include <sys/capability.h>
> #include <unistd.h>
> #include <stdio.h>
> int main() {
>     cap_set_proc(cap_from_text("all-eip")); /* drop all caps */
>     setuid(getuid()); /* drop root. this call succeeds */
>     setuid(0); /* this should fail! but doesn't */
>     printf("%d\n", getuid()); /* we regained root. prints 0 */
>     return 0;
> }
>
> (If we don't use capabilities at all, and take out the cap_set_proc
> line, then the program behaves as expected, and doesn't allow us to
> regain root.)
```

Yup, that does seem weird, indeed. The semantics of uid syscalls is already weird, and when you introduce POSIX capabilities, they get even weirder. What is going on here is that `setuid(getuid())` will only modify the saved uid if `CAP_SETUID` is enabled. Normally, `CAP_SETUID` is enabled if and only if the effective uid is zero. However, dropping all capabilities changes this, and consequently your attempt to drop root still leaves `suid==0`, which is what enables the subsequent `setuid(0)` to succeed. You can confirm this by reading the source code, which can be found in `kernel/sys.c:sys_setuid()`.

Nonetheless, even though I can explain what the OS is doing, I can't explain why it is doing that. I have no idea why someone decided that `setuid()` should leave the saved uid unmodified if `CAP_SETUID` is not present.

Re: Why can setuid programs regain root after dropping it when using capabilities?

My recommendation: Use `sys_setresuid()`. It has by far the most intuitive semantics, and POSIX capabilities don't trigger any surprising modifications to its behavior.

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to `majordomo@xxxxxxxxxxxxxxxxx`

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

- Prev by Date: ***Re: MPlayer broken under 2.6.15-rc7-rt1?***
- Next by Date: ***Re: [PATCH 1 of 3] Introduce memcpy to io32***
- Previous by thread: ***2.6.14 on sparc64 – ext3: journal block not found***
- Next by thread: ***2.6.14.5: segfault / oops with ide-scsi***
- Index(es):
 - ◆ ***Date***
 - ◆ ***Thread***